

CLAIMS:

1. A method of scheduling a plurality of tasks in a data processing system having a plurality of processors, comprising the steps of:

for each task of the plurality, providing suspension data specifying suspension of the task based on memory used thereby;

allocating each of said plurality of tasks to a processor of said plurality of processors; and

each processor performing the steps of:

- (i) processing one the tasks allocated to the processor,
- (ii) monitoring for an input indicative of memory used by the task matching the suspension data associated with the task,
- (iii) suspending said task on the basis of said monitored input, and
- (iv) processing a different one of the tasks allocated to the processor.

2. The method of claim 1, wherein:

allocation of a task to a processor is based on one of -

- a. fixed allocation having every task allocated to a particular processor;
- b. variable allocation having every task allocated to every processor; and
- c. mixed allocation having every task allocated to a subset of said plurality of

processors.

3. The method of claim 2, wherein said input comprises data indicative of a suspension request.

4. The method of claim 3, further comprising the steps of:

receiving first data identifying maximum memory usage associated with the plurality of tasks;

receiving second data identifying memory available for processing the plurality of tasks;

identifying, on the basis of the first and second data, whether there is sufficient memory available to process the tasks; and

wherein, said monitoring and suspending steps are performed only in response to identifying insufficient memory.

5. The method of claim 4, further comprising the steps of:
monitoring termination of tasks; and
in response to a task termination, repeating said step of identifying availability of memory in response to a task terminating.

6. The method of claim 5, in which, in response to identifying sufficient memory to execute the remaining tasks, the monitoring step is deemed unnecessary.

7. The method of claim 6, further comprising the steps of:
receiving first data identifying maximum memory usage associated with the plurality of tasks;
receiving second data identifying memory available for processing the plurality of tasks;
identifying, on the basis of the first and second data, whether there is sufficient memory available to process the tasks; and
wherein, said monitoring and suspending steps are performed only in response to identifying insufficient memory.

8. The method of claim 7, further comprising the steps of:
monitoring termination of tasks; and
in response to a task termination, repeating said step of identifying availability of memory in response to a task terminating.

9. The method of claim 8, in which, in response to identifying sufficient memory to execute the remaining tasks, the monitoring step is deemed unnecessary.

10. The method of claim 2, further comprising the steps of:
receiving first data identifying maximum memory usage associated with the plurality of tasks;

receiving second data identifying memory available for processing the plurality of tasks;

identifying, on the basis of the first and second data, whether there is sufficient memory available to process the tasks; and

wherein, said monitoring and suspending steps are performed only in response to identifying insufficient memory.

11. The method claim 10, further comprising the steps of:

monitoring termination of tasks; and

in response to a task termination, repeating said step of identifying availability of memory in response to a task terminating.

12. The method according to claim 11, in which, in response to identifying sufficient memory to execute the remaining tasks, the monitoring step is deemed unnecessary.

13. A method of scheduling a plurality of tasks in a data processing system having a plurality of processors, comprising the steps of:

for each task of the plurality, providing suspension data specifying suspension of the task based on memory used thereby;

allocating each of said plurality of tasks to a processor of said plurality of processors based on one of -

- fixed allocation having every task allocated to a particular processor,
- variable allocation having every task allocated to every processor, and
- mixed allocation having every task allocated to a subset of said plurality of

processors; and

each processor performing the steps of:

- (i) processing one the tasks allocated to the processor,
- (ii) monitoring for an input indicative of memory used by the task matching the suspension data associated with the task,
- (iii) suspending said task on the basis of said monitored input, and
- (iv) processing a different one of the tasks allocated to the processor.

14. A scheduler for use in a data processing system having a plurality of processors, the data processing system being arranged to execute a plurality of tasks on said plurality of processors and having access to a specified amount of memory for use in executing the tasks, the scheduler comprising:

a data receiver arranged to receive data identifying maximum memory usage associated with a task;

an evaluator arranged to identify, on the basis of the received data, whether there is sufficient memory to execute the tasks;

an allocator arranged to allocate each of said plurality of tasks to a processor of said plurality of processors based on one of -

- a. fixed allocation having every task allocated to a particular processor,
- b. variable allocation having every task allocated to every processor, and
- c. mixed allocation having every task allocated to a subset of said plurality of processors;

a selector arranged to select at least one task for suspension during execution of the task, said suspension coinciding with a specified memory usage by the task; and

a scheduler 501 arranged to initiate execution of said allocated task and suspend said selected task,

wherein, in response to the evaluator identifying that there is insufficient memory to execute the plurality of tasks, the selector selects at least one task for suspension, on the basis of its specified memory usage, and the specified amount of memory available to the data processing system, and the scheduler suspends execution of the at least one selected task in response to the task using the specified memory.

15. A scheduler according to claim 14, wherein the evaluator is arranged to monitor termination of tasks, and in response to a task terminating, to identify whether there is sufficient memory to execute the remaining tasks.

16. A scheduler according to claim 15, wherein, in response to the evaluator identifying sufficient memory to execute the remaining tasks, the selector is arranged to deselect said selected at least one task.

17. A data processing system having a plurality of processors arranged to execute a plurality of tasks, the data processing system including:

memory arranged to hold instructions and data during execution of a task;

receiving means arranged to receive data identifying maximum memory usage associated with a task and data specifying preemptability of the task;

evaluating means arranged to identify, on the basis of the received data, whether there is sufficient memory to execute the tasks;

an allocator arranged to allocate each of said plurality of tasks to a processor of said plurality of processors based on one of -

- fixed allocation having every task allocated to a particular processor,
- variable allocation having every task allocated to every processor, and
- mixed allocation having every task allocated to a subset of said plurality of

processors; and

a scheduler arranged to schedule execution of the tasks on the basis of input received from the evaluating means,

wherein, in response to identification of insufficient memory to execute the plurality of tasks, the scheduler is arranged to suspend execution of at least one task in dependence on memory usage by the task.

18. A method of transmitting data to a data processing system having a plurality of processors, the method comprising:

allocating each task of a plurality of tasks to a processor of said plurality of processors based on one of -

- fixed allocation having every task allocated to a particular processor,
- variable allocation having every task allocated to every processor, and
- mixed allocation having every task allocated to a subset of said plurality of

processors;

transmitting data for use by the data processing system in processing each task of said plurality of tasks;

transmitting suspension data specifying suspension of each task of said plurality of tasks based on memory usage during processing thereof;

wherein, the data processing system is configured to perform a process comprising:

monitoring for an input indicative of memory usage of each task matching the suspension data associated with the task; and
suspending processing of said each task on the basis of said monitored input.

19. A method according to claim 18, wherein the suspension data includes data identifying maximum memory usage associated with said each task.

20. A method according to claim 18, wherein the suspension data identifies at least one point at which processing of each task can be suspended, based on memory usage of said each task.

21. A method according to claim 20, wherein said each task comprises a plurality of sub-jobs and said data identifying at least one point at which processing of said each task can be suspended corresponds to each such sub-job.

22. A method according to claim 20, wherein the suspension data includes data identifying maximum memory usage associated with said each task.

23. A method according to claim 22, wherein said each task comprises a plurality of sub-jobs and said data identifying at least one point at which processing of said each task can be suspended corresponds to each such sub-job.

24. A method of configuring a plurality of tasks for use in a data processing system having a plurality of processors, the method including associating suspension data with the task, the suspension data specifying suspension of the task based on memory usage associated therewith, wherein the data processing system is arranged to perform a process in respect of a plurality of tasks executing on a plurality of processors, the process comprising:

allocating each task of said plurality of tasks to a processor of said plurality of processors based on one of -

- fixed allocation having every task allocated to a particular processor,
- variable allocation having every task allocated to every processor, and

- mixed allocation having every task allocated to a subset of said plurality of processors;
monitoring for an input indicative of memory usage of the task matching the suspension data associated with the task; and
suspending processing of said task on the basis of said monitored input.

25. A computer program, comprising a set of instructions arranged to cause a processing system to perform the method according to of claim 1.

26. A computer program, comprising a set of instructions arranged to cause a processing system to perform the method according to of claim 2.